

METHOD FOR DESIGN OF NETWORKS BASED ON p -CYCLES

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims priority from US provisional application no. 60/430,931 filed December 5, 2002.

BACKGROUND OF THE INVENTION

01 p -Cycles offer a promising new approach to optical network survivability, summed up by the notion of “ring like speed with mesh-like efficiency”. Being based on closed cyclic paths of protection capacity, p -cycles offer ring-like speed and pre-configured simplicity but are essentially as efficient as span-restorable mesh networks thereby offering three to six times greater demand-carrying capability than rings for a given transmission capacity. Unlike rings, p -cycles protect straddling failures as well as on-cycle failures and allow working paths to take shortest routes. This combination of properties, suggest the prospect of an optical network that is survivable to any single span failure with as little as ~35% redundancy, depending on graph topology and demand pattern. In contrast, optical rings and fiber-level cycle double covers are *at best* 100% redundant (and often much higher) in terms of spare and unused working capacity. p -Cycles also provide the much-touted “50 ms” speed, because only two nodes do any switching, and the failure-dependent local switchover actions are BLSR-like and known in advance.

02 p -Cycles are like rings but with support for the protection of *straddling* span failures as well as the usual ring-protection of spans of the ring itself (*on-cycle* failures). A straddling span has its end-nodes on the p -cycle, but is not itself part of the p -cycle, like a chord on a circle. With p -cycles, working paths also take any desired route over the graph and are not constrained to follow ring routings. When an on-cycle span fails, the surviving arc of the cycle is used just as in a BLSR ring. However, the same p -cycle is also accessible to support restoration of a straddling span failure, in which case two restoration paths are available from each unit of p -cycle protection capacity. In the limit of a full set of straddling span relationships, an N -hop p -cycle can protect up to $N(N-2)$ units of working capacity, making it up to $N-2$ times more efficient than a corresponding ring. However, the design of a min-cost set of p -cycles to protect a given set of

working flows is an NP-hard problem. The basic formulation generates large problem files that can be difficult to solve to optimality, primarily because of the size of the set of candidate cycles to consider. This is especially true when the jointly optimized problem or the consistent wavelength assignment problem is attempted. Some investigators are pursuing pure heuristics for the problem, and a fully distributed p -cycle forming process is known. This invention seeks to provide an improved method for designing a telecommunications network that optimizes routing of working demands and spare capacity, that is, that provides a solution to what is known as the joint optimization problem.

SUMMARY OF THE INVENTION

03 According to an aspect of this invention, a method of designing a telecommunications network including the method step of pre-selecting a reduced number of “high merit” p -cycle candidates which are then provided to an otherwise unchanged optimal solution model is presented.

04 Thus, there is provided a method of providing a mesh telecommunications network with spare capacity arranged in pre-configured cycles by pre-selecting a set of candidate cycles for forming into pre-configured cycles and allocating working paths and spare capacity in the mesh network based on the set of candidate cycles. The network is then provided with the spare capacity in pre-configured cycles determined by the allocation. The allocation of working paths and spare capacity may be carried out by joint optimization.

05 Pre-selecting candidate cycles may include ranking a set of closed paths in the mesh telecommunications network according to the degree to which each closed path protects spans on and off the closed path, and selecting candidate cycles from the set of closed paths, and preferably takes into account the cost of the closed path.

06 According to a further aspect of the invention, pre-selecting candidate cycles comprises:
a) determining a topological score of the closed paths in the set of closed paths, where the topological score of said closed path is increased by a value for each span within said closed path

that is protected by said closed path, and increased by a larger value for each span not on said closed path that is protected by said closed path;

b) determining the a priori efficiency of each closed path, where the a priori efficiency of a closed path is determined by taking the ratio of the topological score of said closed path with the cost of said closed path; and

c) choosing a select number of closed paths based on the a priori efficiency to be the pre-selected candidate cycles.

BRIEF DESCRIPTION OF THE DRAWINGS

07 There will now be given a brief description of preferred embodiments of the invention, by reference to the drawings, by way of illustration only and not limiting the scope of the invention, and in which:

Fig. 1(a) depicts a network with a p -cycle that protects a working path which is following its shortest route over the graph.

Fig. 1(b) depicts a network with a p -cycle the choice of which is also coordinated with a changed working path routing so that the total of working and spare capacity required is minimized. Fig. 2(a) is a flow chart representing the steps in a non-joint approach to network design.

Fig. 2(b) is a flow chart representing the steps in a joint approach to network design.

Fig. 3(a) and (b) are examples of calculating the topological score of a p -cycle according to an aspect of the invention.

Fig. 4(a) and (b) are examples of calculating the “a priori efficiency” of a p -cycle according to an aspect of the invention.

Fig. 5 depicts a network of European cities used as an example of the invention.

Fig. 6 compares results from using the selection criteria presented in this disclosure with those from a random selection in the non-joint case based on the COST 239 network.

Fig. 7 compares results from using the selection criteria presented in this disclosure with those from a random selection in the joint case based on the COST 239 network.

Fig. 8 compares results from using the selection criteria presented in this disclosure with those from a random selection in the non-joint case based on an additional test network.

Fig. 9 compares some details of the joint and non-joint designs in the COST 239 network.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

08 In this patent document, a mesh telecommunications network (also often called a “transport” network or “optical network”) is a telecommunications network formed from plural nodes connected by plural spans. In this patent document, the word “comprising” is used in its non-limiting sense to mean that items following the word in the sentence are included and that items not specifically mentioned are not excluded. The use of the indefinite article “a” in the claims before an element means that one of the elements is specified, but does not specifically exclude others of the elements being present, unless the context clearly requires that there be one and only one of the elements.

09 This disclosure addresses two open and inter-related issues about the design of p -cycle based networks. One of these is to reduce the complexity of solving optimal p -cycle design problems, making it practical to continually re-optimize a p -cycle based network in service, adapting to changing demand patterns, or practical to do many different design and planning studies in a reasonable time using standard computer planning software for network design. The second advance is a first research use of the above solution technique to study how the efficiency of a p -cycle network increases under *joint* optimization of the working path routes with p -cycle placement.

10 The aspect of *jointness* in a p -cycle design problem will now be discussed. The issue is that one can either first route the working demands via shortest paths (or any other means) and then solve a corresponding minimum *spare* capacity allocation problem (the non-joint problem), or, attempt to optimize the choice of working routes in conjunction with the placement of spare capacity together (i.e., jointly) to minimize *total* capacity (the joint problem). An example of the effect of solving the joint problem is shown in Figs. 1(a) and (b). Network 10 is shown with flow 12 representing 2 units of capacity 2λ , and p -cycle 14 consisting of spans 101-106. In Fig. 1(a), flow 12 passes through spans 101 and 102, while in Fig. 1(b), flow 16 passes through spans 107 and 108. In Fig. 1(a), p -cycle 14 requires 2 units of spare capacity on spans 101-106, while Fig. 1(b) reduces the units of space capacity required for p -cycle 14 to one unit on spans 101-106

by increasing the length of the path of the working flow from 2 to $2+\epsilon$. Together these figures show how potentially advantageous it can be in terms of total capacity reduction, if the right choices of working path routing can be made, in conjunction with p -cycle placement choices, so as to realize the kind of savings evident in the example of Figure 1.

11 A flow chart showing an example of how the non-joint problem is solved is shown in Fig. 2(a), where the first step in 202 is to route all lightpath requirements via the shortest path. Step 204 follows from this in defining all working paths and working capacities. To solve for the required p -cycles and spare capacity in 212, one can either enumerate all graph cycles as possible p -cycles in 206 and then use an ILP solution for p -cycle formation in 208, or one can use another heuristic algorithm(s) for p -cycle formation in 210. This process can be contrasted with the flow chart in Fig. 2(b), where the first steps 220 and 222 are to enumerate eligible working routes and enumerate graph cycles as potential p -cycles, respectively. These two sets are combined and solved in a single ILP solution in 224, from which the p -cycles and spare capacity are defined in 226, and the working routes and capacity are defined in 228.

12 It is known by published work that joint optimization reduces total capacity by as much as 28%, in conventional span restoration but under 10% for path restoration. But the corresponding benefit of joint design is not yet known for p -cycles, in part because of the additional complexity of solving the joint problem formulation. The pre-selection heuristic presented in this disclosure has been applied to obtain practical and virtually optimal solutions to the joint p -cycle design problem.

13 As it is desired to reduce the complexity of solving optimal p -cycle design problems, it is necessary to define the criteria used in pre-selecting potential p -cycles. Two pre-selection metrics, *topological score* and *a priori efficiency* are based on insights about what makes for the most efficient p -cycles in the context of a given network design. The *topological score* (TS) is defined by equation 1:

$$TS(j) \triangleq \sum_{i \in S} x_{ij} \quad (1)$$

where S is the set of spans, $x_{ij}=1$ if span i is part of cycle j , $x_{ij}=2$ if span i straddles cycle j and

$x_{ij}=0$ otherwise. Figs. 3(a) and (b) show examples of how the TS is calculated for two different p -cycles 32 and 34 respectively for the network 30. Cycle 32 is composed of 6 spans, with no straddling spans, while cycle 34 is composed of 7 spans, with two straddling spans. For each on-cycle span protection relationship, the TS is increased by 1, such as for spans 300 to 306 in Fig. 3(a) and spans 305 to 311 in Fig. 3(b). For each straddling span protection relationship, the TS is increased by 2, such as for spans 312 and 313 in Fig. 3(b). Thus, $TS(32)$ in Fig. 3(a) receives a score of 6 because it protects 6 on-cycle spans with no straddling span protection, while $TS(34)$ in Fig. 3(b) receives a score of 11 because it protects 7 on-cycle spans with 2 straddling spans being protected.

14 The *a priori efficiency* (AE) of a cycle j is defined by equation 2:

$$AE(j) \triangleq TS(j) / \sum_{(i \in S | x_{ij}=1)} C_i \quad (2)$$

where $TS(j)$ is the topological score for cycle j calculated previously and C_i is the cost or distance of span i . An example of calculating the AE of two cycles is shown in Figs. 4(a) and (b) where the network 30 and the p -cycles 32 and 34 of Figs. 3(a) and (b) are depicted. The cost C_{32} for cycle 32 is 6 and the cost C_{34} is 7, being 6 and 7 "hops" in length, respectively. Thus, using $TS(32)$ and $TS(34)$ calculated previously, we find that $AE(32) = 1$ and $AE(34) = 1.57$. Note that rings always have $AE = 1$.

15 Whereas the basic solution model requires representation of all distinct cycles (possibly up to a circumference limit), we simply rank the set of all distinct cycles by either TS or AE measures, and use only a limited number of the top-ranked candidates for representation in the optimal solution model. In the example above, cycle 34 would be ranked higher than cycle 32.

16 Once the set of candidate cycles of the network graph have been characterized in this way, the problem can be solved using, for example, an integer linear programming (ILP) formulation, where the objective function minimizes the total cost of spare capacity and (for the joint problem) working capacity. ILP formulations are well known in the art and need not be further described here. This function is subject to:

A. All lightpath requirements are routed.

B. Enough WDM channels (or working channels in general) are provided to accommodate the routing of lightpaths in A.

C. The selected set of p -cycles give 100% span protection.

D. Enough spare channels are provided to create the p -cycles needed in C.

E. Integer p -cycles decision variables and integer capacity.

17 Applying the pre-selection criteria can be particularly useful in the joint optimization problem, where the formulation generates large problem files that can be difficult to solve optimally if there is no pre-selection.

18 For a test case study, the COST 239 test case of 11 major European cities inter-connected by 26 spans and exchanging a total of 176 lightpath demands in a random mesh-like pattern shown in Fig. 5, varying from 1 to 11 lightpaths per node pair (average of 3.2) was considered, following Schupke in “Optimal Configuration of p -Cycles in WDM Networks,” D.A. Schupke, C.G. Gruber, A. Autenrieth, *ICC 2002*. Selected results are summarized in Figs. 6, 7, and 9, where, in Figs. 6 and 7, the horizontal axis is the number of candidate cycles used in the exercise, and the vertical axis is the total capacity optimality gap in percentage. Figure 6 shows results comparing a random set of selected candidate cycles with those selected using the selection criteria presented in this disclosure for the non-joint design case. Note that using the criteria allows us to reach near 0% gap in optimization using 500 candidate cycles, while randomly selecting candidate cycles reaches a 0% gap around 2000. (And a factor of four in size of the problem representation can mean much more than four times longer run-times). Figure 7 shows results comparing a random set of selected candidate cycles with those selected using the selection criteria presented in this disclosure for the joint design case. Using the selection criteria allows us to reach ~0% gap using 200 candidate cycles, while the random case still requires 2000 candidates. Figure 9 shows a bar graph of the capacity of the joint and non-joint design cases, with the capacity on the vertical axis. Note that the joint case uses 5% more working capacity, due to longer paths being chosen in some cases, but it also uses 43% less spare capacity giving a total network capacity reduction of 13%, and an overall network redundancy of 39%. As Figure 9 also shows, a very significant extra practical advantage of the joint design is that it requires only 4 distinct p -cycle systems to be built, whereas the non-joint design requires 7 distinct p -

cycles. Although hard to quantify, such reduction of “systems count” when designing transport networks is widely known to be of significant capital and operational savings value over and above the direct value of using less total capacity.

19 Note that when given only the 250 top-ranked cycles to consider by the AE metric the non-joint problem was solved to within 1.14% of optimal 857 times faster than the unrestricted non-joint problem. The joint problem is solved to within 0.16% of optimal in 1/17th of the time if provided with only 50 cycles as pre-selected by the ranked AE criterion.

20 Results from another test network are shown in Fig. 8. The network used consisted of 20 nodes, 40 spans, and 190 demand pairs. Compared are the results of a random selection for the joint design case (JCA) and using the AE selection criteria. While more candidate cycles were required for this more complex network, the selection criteria allows the network to drop below a 0.5% gap much sooner than the random case.

21 Two points concerning the impact and relevance of the invention can be drawn from the preceding test case. Firstly, we note that, at least in COST 239, the joint p -cycle design is as efficient as previously studied dynamic path-restorable designs in other studies to date. Such high efficiency is a direct benefit in terms of reduced cost or greater revenue from the same facilities, but an efficient network is also inherently a more *flexible* network because less of its resources are tied up for protection. Secondly, the simple process of pre-selecting candidate p -cycles by the AE metric greatly reduces p -cycle solution times so much that it may be practical to continually re-compute the optimal p -cycle configuration on-line as the network demand pattern evolves. This helps greatly to remove some prior objections to the practicality of p -cycle based networks and enables the vision of a continually adapting background layer of p -cycles.

22 It is possible have “too elite” a population of candidate cycles and it may be desirable to dilute the population with a few other types of candidates. The basic framework is one within which many specific heuristic ideas can be tried, all having to do with defining the reduced set of elite cycles to consider. First, some experience with memory and run times may show, for

example, that a budget of 10,000 cycles is realistic to work with. The budget can be used up representing any number of mixed strategies for populating the elite P set.

An example could be:

Admit the 5,000 cycles found by the score-based selection above.

Add the 2,000 cycles with most absolute number of straddlers.

Add the 2,000 of the longest cycles.

Add 1,000 random cycles.

By itself the first set of cycles may not necessarily ensure feasibility. When choosing only individually elite cycles, there is no strict guarantee that a cycle will be represented that would cover, for example, a very long degree-2 chain connected to an otherwise highly connected mesh. However, cycles in batch three above definitely cover that eventuality.

23 Immaterial modifications may be made to the embodiments described here without departing from the invention.